

Claims

[c1] 1. A data management system, said system characterized as a composite system, the system comprising a plurality of processes; each process having an interface and implementing at least one respective service defined by that interface; a first invocation of the at least one respective service by a transaction resulting in the creation of a first transaction local to the process thereof, the first local transaction being a child of the invoking transaction and being parent of any transaction triggered by invocation of a service of another process; a second invocation of the at least one respective service by a transaction resulting in the creation of a second transaction local to the process thereof, the first local transaction being a child of the invoking transaction and being parent of any transaction triggered by invocation of a service of another process; each transaction comprising the local invocation of at least two resources, each such invocation leading to an internal participant instance for globalcommit; each such resource and server invocation also containing application-level comments regarding the nature of the

invocation, such comments comprising a portion of any globalcommit message exchanges; each such comment being presented to a human administrator in the event of a global commitment failure; each process characterized in that if the first transaction and the second transaction conflict but are both children of a same invoking transaction, then the first transaction and the second transaction and ancestors, if any, of the first transaction and the second transaction are not executed concurrently; each process further characterized in that each transaction local thereto is independently handled at the process; each process making scheduling and recovery decisions independent of any centralized component.

- [c2] 2. The system of claim 1 wherein each transaction defining a start associated therewith, and a commit associated therewith; and wherein the condition of not executing concurrently comprises the condition that the start associated with the second transaction must happen after the commit associated with the first transaction.
- [c3] 3. The system of claim 1 each server resource further comprising a lock table; each invocation of an internal resource including an

identifier of the invoking transaction in an entry in the lock table associated with the invocation; each invocation of an internal resource comprising checking the invocation against any entries in the lock table to determine whether the invocation has an identifier of the invoking transaction matching an identifier of an entry in the lock table, a match defining the condition of the invoking transaction being a child of the same invoking transaction as the invoking transaction of the entry giving rise to a match.

- [c4] 4. The system of claim 1 wherein each resource invocation further comprises an undo operation, the undo operation being local to the resource of the service.
- [c5] 5. The system of claim 4 wherein in the case of aborting transactions that are children of the same invoking transaction, all undo transactions for the same invoking transaction are executed on each resource in the reverse order of their respective executions.
- [c6] 6. A data management system, said system characterized as a composite system, the system comprising a plurality of processes; each process having an interface and implementing at least one respective service defined by that interface; a first invocation of the at least one respective service by

a transaction resulting in the creation of a first transaction local to the process thereof, the first local transaction being a child of the invoking transaction and being parent of any transaction triggered by invocation of a service of another process;

a second invocation of the at least one respective service by a transaction resulting in the creation of a second transaction local to the process thereof, the first local transaction being a child of the invoking transaction and being parent of any transaction triggered by invocation of a service of another process;

each process characterized in that if the first transaction and the second transaction conflict but are both children of a same invoking transaction, then the first transaction and the second transaction are not executed concurrently;

each process further characterized in that each transaction local thereto is independently handled at the process;

each process making scheduling and recovery decisions independent of any centralized component.

[c7] 7. The system of claim 6 wherein the root transaction is able to dynamically set concurrency preferences for the resulting distributed transaction, based on client needs.

[c8] 8. A data management system, said system characterized as a composite system, the system comprising a plurality of processes; each process having an interface and implementing at least one respective service defined by that interface; invocation of the at least one respective service by a strand of the invoking transaction and being parent of any transaction triggered by invocation of a service of another process; each process further characterized in that each transaction local thereto is independently handled at the process; each process making scheduling and recovery decisions independent of any centralized component. triggered by invocation of a service of another process, each process further characterized in that each transaction local thereto is independently handled at the process, each process making scheduling and recovery decisions independent of any centralized component, the method comprising the steps of: propagating from a first process to a second process a message indicative of a globalCommit operation with respect to a root transaction, said message also indicative of a number or identifying list of invocations which the first process has made to the second process on behalf of the root transaction;

within the second process, comparing the number or list indicated in the message with a count or list within the second process of the number or list of invocations which have been made on behalf of the root transaction; in the event the comparison yields a non-match, aborting the transaction.

- [c9] 9. The system of claim 8 wherein each process is built using Java.
- [c10] 10. A method for use with a data management system, said system characterized as a composite system, the system comprising a plurality of processes, each process having an interface and implementing at least one respective service defined by that interface, invocation of the at least one respective service by a transaction resulting in the creation of a transaction local to the process thereof, the local transaction being a child of the invoking transaction and being parent of any transaction triggered by invocation of a service of another process, each process further characterized in that each transaction local thereto is independently handled at the process, each process making scheduling and recovery decisions independent of any centralized component, the method comprising the steps of:
propagating from a first process to a second process a message indicative of a globalCommit operation with re-

spect to a root transaction, said message also indicative of a number or list of invocations which the first process has made to the second process on behalf of the root transaction;

within the second process, comparing the number or list indicated in the message with a count or list within the second process of the number or list of invocations which have been made on behalf of the root transaction; in the event the comparison yields a match, proceeding with the globalCommit operation.

- [c11] 11. A method for use with a data management system, said system characterized as a composite system, the system comprising a plurality of processes, each process having an interface and implementing at least one respective service defined by that interface, invocation of the at least one respective service by a transaction resulting in the creation of a transaction local to the process thereof, the local transaction being a child of the invoking transaction and being parent of any transaction triggered by invocation of a service of another process, each process further characterized in that each transaction local thereto is independently handled at the process, each process making scheduling and recovery decisions independent of any centralized component, the method comprising the steps of:

propagating from a first process to a second process a message indicative of a globalCommit operation with respect to a root transaction, said message also indicative of a number or list of invocations which the first process has made to the second process on behalf of the root transaction;

within the second process, comparing the number or list indicated in the message with a count or list within the second process of the number or list of invocations which have been made on behalf of the root transaction; in the event the comparison yields a non-match, aborting the transaction.

[c12] - 2 -

12. A distributed system, said system characterized as a composite system, the system comprising a plurality of processes; each process having an interface and implementing at least one respective service defined by that interface; each or any globalCommit message exchange between processes also carrying information about the actual work being committed (or being agreed upon).

[c13] 13. The system of claim 12, such information being logged for recoverability in the event of a crash, such information being used for assistance (or presented) at any time before, during or after global commitment.

- [c14] 14.The system of claim 12 or 13, wherein any globalCommit requires a registration, and wherein the registration for a globalCommit also carries information about the actual work being committed.
- [c15] 15.A method for use in a distributed system, said system characterized as a composite system, the system comprising a plurality of processes, each process having an interface and implementing at least one respective service defined by that interface, the method comprising the step of: for each globalCommit message exchanged between processes, including also information about the actual work being committed (or being agreed upon).
- [c16] 16.The method of claim 15 further comprising the step of logging such information for recoverability in the event of a crash, such information being used for assistance (or presented) at any time before, during or after global commitment.
- [c17] 17.The method of claim 15 or 16 further comprising the step of propagating a registration for a globalCommit, wherein the registration for a globalCommit also carries information about the actual work being committed.
- [c18] 18.A distributed system, said system characterized as a composite system, the system comprising a plurality of

processes; each process having an interface and implementing at least one respective service defined by that interface; wherein the root invocation (transaction) or, alternatively, the root's human user is allowed to dynamically set its/his concurrency preferences for the entire invocation.

- [c19] 19. The system of claim 18 wherein the root invocation (transaction) propagates the concurrency preferences with each or any child invocation it makes.
- [c20] 20. The system of claim 19 wherein each invocation propagates the concurrency preferences as it has received them from the root invocation.
- [c21] 21. The system of claim 20 wherein the propagated concurrency preferences (at any level in the root invocation's invocation hierarchy) specify the extent to which shared resource access is desired or allowed or denied among descendant invocations of the root invocation (or user) and other, concurrent invocations who are also descendants of the same root.
- [c22] 22. The system of claim 20 wherein the propagated concurrency preferences (at any level in the root invocation's invocation hierarchy) specify the extent to which shared resource access is desired or allowed or denied among

descendant invocations of the root invocation (or user) and other, concurrent invocations who are not descendants of the same root.

[c23] 23. A data management system, referred to as service, comprising:

- One or more operations that can be invoked by remote clients;
- Some or all such remote clients having one or more associated contexts or transaction contexts;
- An invocation by a remote client also containing partial or complete information indicating or containing said client's context or contexts;
- An invocation, by a remote client, of an operation leading to a new transaction different from, but possibly related to, any existing client transaction;
- Such an operation-level transaction being committed before the client context is terminated (before global-Commit notification);
- The service maintaining an undo operation for such a committed operation;
- A failing or failed remote client context leading to the execution of the undo operations of the corresponding committed invocations in the service.

[c24] 24. The system of claim 23 where some or all undo operations are executed in an order that is the reverse of the

order of their original counterparts.

- [c25] 25.The system of claim 23 where in addition the undo operations are chosen or defined in the same system as the one where the corresponding normal operations were executed.
- [c26] 26.The system of claim 23 where some or all undo operations are unknown to a remote client or its context.
- [c27] 27.The system of claim 23 where some or all undo operations are executed after a timeout and independent of whether the client's context outcome requires such undo.
- [c28] 28.The system of claim 23 where an undo operation's effects are confined to the data managed by the service on which the undo operation is maintained, even if the original operation involved other services.
- [c29] 29.The system of claim 23 where the service keeps locks to ensure that undo operations can be executed correctly.
- [c30] 30.The system of claim 23 where client context-related information is also part of any global commit message exchanges.
- [c31] 31.The system of claim 23 where client context informa-

tion includes application-specific data.

- [c32] 32.The system of claim 31 where client all or part of the context information is logged, i.e. stored on persistent storage, and retrievable by a human administrator.
- [c33] 33.The system of claim 23 where the service accepts messages indicative of which previously committed operations have to be undone.
- [c34] 34.The system of claim 23 where the service accepts messages indicative of which previously committed operations do not have to be undone.
- [c35] 35.The system of claim1 where some or all invocations are message-based or asynchronous.
- [c36] 36.The system of claim 23 where some or all invocations are synchronous.
- [c37] The system of claim 23 where the client can request the undo executions of its invocations at the service while still allowing globalCommit in the end.